

AUTOMATED PROGRAM RESOURCE IDENTIFICATION AND ASSOCIATION

FIELD OF THE INVENTION

This invention is directed to the field of computer security. It is more particularly directed to the detection of authorization usage for computer programs.

BACKGROUND OF THE INVENTION

Java 2 has a security architecture that is intended to protect client and server systems from malicious code that has been dynamically installed (e.g., mobile code). For example, applet code can be downloaded from the Internet into a Web browser, or uploaded via RMI into a server application. The Java 2 security system includes an authentication subsystem and an authorization subsystem. In the discussion that follows, we will concern ourselves only with the authorization subsystem.

Java 2 authorization is based on a set of ProtectionDomains, where each Java application class is loaded via a SecureClassLoader into the Java virtual machine. In the reference implementation of the Java virtual machine, classes part of the Java runtime are loaded from the "boot classpath", by a special class loader sometimes referred to as the "primordial class loader." These classes are considered fully trusted and are implicitly given AllPermission. Each application class refers to a ProtectionDomain. This ProtectionDomain object includes a CodeSource and a set of Permissions. The CodeSource is authentication information (the code base URL, and the digital certificates used to sign the code if the code was signed). The Permissions are a set of

1 Permission objects representing the permissions, or rights, the code has with respect to privileged
2 operations. It is the computation of this set of Permission objects that is the subject of this
3 invention.

4 Prior to deploying (application or library) code in Java, an important question arises: "What Java
5 Permission(s) are required to enable the code to run in a system which has the security
6 authorization subsystem enabled?" Currently, this problem is solved empirically. The code
7 writer reads the documentation for any libraries (including the Java runtime libraries) invoked by
8 the code and determines whether any of these libraries requires that the code have specific
9 Permission(s) in order to run. Unfortunately, this documentation is often missing information on
10 used Permission(s), or the documentation is out of date. In the absence of reliable
11 documentation, the developer runs the new code and waits until a SecurityException is thrown
12 due to an insufficient set of Permission(s) allocated to the new code. The developer then adds
13 any additional used Permission(s) to the set granted to the code under test. This trial-and-error
14 process is repeated until no more SecurityException failures occur.

15 An analogous scenario can be depicted on the system where mobile code is dynamically
16 installed. When mobile code is downloaded on a system, the system administrator (e.g., the web
17 browser user) has to figure out what Permission(s) that code needs in order for it to work without
18 causing SecurityException failures. The system administrator should therefore rely on what the
19 code developer/distributor recommends, with the risk that too many Permission(s) are granted
20 and security holes are opened. Alternatively, the system administrator should first run the code
21 with no Permission(s), look at the SecurityExceptions thrown, and incrementally grant
22 Permission(s) as is necessary. This process can be tedious and error-prone. In addition,
23 insufficient code testing may result in a wrong set of Permission(s) being granted, which may
24 cause security exposures or code instability.

1 Rather than relying on empirical, or ad hoc, methods of determining what Permission(s) are used
2 to run application/library code, we describe a process by which the used Permission(s) for a
3 software component (e.g., library, application, etc.) can be computed.

4 **Java 2 Authorization Algorithm**

5 When the Java 2 authorization subsystem is called, a test is made to determine whether all of the
6 methods in the runtime stack are authorized to perform the privileged operation. Specifically,
7 when a Java program is running, the Java virtual machine builds a call stack, with each element
8 of the stack being an activation record that holds state information about a method that has been
9 called during program execution. Each activation record holds a reference to the called method,
10 as well as local variables, etc. From the referenced method, it is possible to determine the
11 method's class. For Java 2 authorization, we are only concerned with the class whose method
12 was called since permissions are associated with classes, not methods. The Java 2 authorization
13 algorithm will test the classes' ProtectionDomains to see if they include a used Permission.

14 As previously noted, each application class loaded via a SecureClassLoader into the Java runtime
15 is associated with a ProtectionDomain object. This object includes Permission object(s)
16 indicating to which resources the class is authorized access. Classes loaded via the boot classpath
17 are implicitly authorized for any privileged operation (the equivalent of AllPermission).

18 The Java 2 authorization algorithm is executed by calling AccessController.checkPermission(*p*),
19 where *p* is a Permission object. The checkPermission() method obtains an
20 AccessControlContext object, which is done by searching the current Java thread to find the
21 activation records for all methods called in the current thread. For each activation record, the
22 associated class for the activated method is identified, and the ProtectionDomain object for that
23 class is obtained. This set of ProtectionDomains for the classes on the thread call stack is called
24 an AccessControlContext. AccessController.checkPermission() determines whether Permission

1 *p*, passed as an argument to the method, is included in each ProtectionDomain in the
2 AccessControlContext. If not, a SecurityException is thrown by the checkPermission() method,
3 indicating that authorization failed. If all ProtectionDomains in the AccessControlContext
4 include the Permission *p*, then checkPermission() silently returns to its caller, indicating that the
5 authorization was successful.

6 The basic authorization algorithm is modified in two significant ways: new threads inheriting an
7 AccessControlContext and calls to AccessController.doPrivileged(). When a new Java Thread is
8 created, the parent thread creates an AccessControlContext and saves it in the child thread as the
9 inherited AccessControlContext. The inherited AccessControlContext includes the set of
10 ProtectionDomain objects from the classes associated with the activation records on the parent
11 thread stack at the time the new (child) thread is created. Subsequently, the child thread call to
12 checkPermission() proceeds as before, but also does an authorization test against the thread's
13 inherited AccessControlContext (if any). That is, all of the ProtectionDomains in the inherited
14 AccessControlContext will also be tested to see if they include Permission *p*. This extra test is
15 done to prevent a child thread from having permissions which are greater than those of the parent
16 thread. If this test were not performed, creating new threads would constitute an easy way to
17 bypass authorization constraints of the parent thread.

18 When a thread calls AccessControlContext.doPrivileged(), a mark is nominally left on the
19 thread's stack at the activation record for the method that called doPrivileged(). When the same
20 thread subsequently needs an AccessControlContext (e.g., to call checkPermission() or for a new
21 Thread), the AccessControlContext that is created only includes ProtectionDomains associated
22 with activation records between, and including, the currently executing method and the activation
23 record for the method that includes the doPrivileged() mark. From a practical perspective, this
24 effectively limits the number of classes requiring the Permission on subsequent calls to
25 checkPermission(). There are variants of AccessControlContext.doPrivileged() which accept an
26 AccessControlContext instance as an argument. As before, a call to checkPermission() will

1 build an AccessControlContext as previously described, but it will also include all
2 ProtectionDomains included in the AccessControlContext passed as an argument to the
3 doPrivileged() call. The effect is to further constrain authorization by requiring additional
4 ProtectionDomains to pass the authorization test.

5 Lastly, many security authorization tests in Java 2 are made through calls to methods in the
6 SecurityManager class. In the reference implementation of this class, most of the authorization
7 tests are made by calling AccessController.checkPermission() with an appropriate Permission
8 object. Additional security policy decisions are implemented by the SecurityManager, but
9 analysis of these policies is outside the scope of this invention. It would also be advantageous to
10 identify the set of authorizations used by a collection of code.

SUMMARY OF THE INVENTION

12 It is therefore an aspect of the present invention to provide methods and apparatus to identify the
13 set of authorizations used by a code collection. To achieve this, we identify paths in a program
14 graph that lead to a node identifying the authorization. We also identify, using a data flow
15 analysis, any necessary parameters passed to methods requesting the authorization. We also
16 traverse the program graph to identify relationships between code collections, for example,
17 thread allocations and thread invocations in programs, to identify additional program paths
18 resulting in authorization tests. Additionally, we aggregate code collections, for example,
19 methods within classes in Java programs, and their associated authorizations, in order to
20 associate associations with the larger, aggregated code collections. Other objects and a better
21 understanding of the invention may be realized by referring to the detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects, features, and advantages of the present invention will become apparent upon further consideration of the following detailed description of the invention when read in conjunction with the drawing figures, in which:

Fig. 1 shows an example of an embodiment of a method employing a computer in accordance with the present invention;

Fig. 2 shows an example of an embodiment of an apparatus of a computing means in accordance with the present invention;

Fig. 3 shows an example of a method employing a computer which enables the identification of a complete set of authorization resources of the collection of code in accordance with the present invention;

Fig. 4 shows an example of an embodiment of an apparatus for a computer providing the advantages of the present invention;

Fig. 5 shows another example embodiment of an apparatus for a computing module in accordance with the present invention;

Fig. 6 shows an example of a flowchart of a method for identifying a complete set of authorization resources of a collection of code, in accordance with the present invention;

Fig. 7 shows another example of an embodiment of a method employing a computer with access to a collection of code in accordance with the present invention; and

Fig. 8 shows another example of an embodiment of a method employing a computer with access to a collection of code; in accordance with the present invention
In an example embodiment of the present invention, the invention includes an apparatus for a computing means. As shown in Figure 2,

DESCRIPTION OF THE INVENTION

The present invention provides methods and apparatus for identifying the set of permissions used by an application class. To achieve this, we identify paths in a program graph that lead to a node whose method is an authorization test. In Java 2, this method includes `AccessController.checkPermission()`. We also identify any authorization object *p* that is passed as an argument to this method. In an advantageous embodiment of this invention the program graph is an invocation graph. Since Java 2 authorization is based on associating `ProtectionDomains` with classes (not methods), a path-insensitive program graph construction algorithm is sufficient. That is, all intra-procedural instruction paths through each method are included in the program graph. In addition, a data flow analysis is performed. In the context of Java 2, the data flow analysis is used to identify the `java.security.Permission` object *p* that is passed as an argument to the `java.security.AccessController.checkPermission()` method. Except as noted below, each class in a path from a root node in the program graph to a node for the `AccessController.checkPermission(p)` method uses `Permission p`. It should be noted that even though the description of the present invention uses the terminology of object oriented programming, the present invention also can be applied to non-object oriented languages and environments.

Call Graph Construction and Characteristics

1 In the advantageous embodiment, we employ a program graph that is path-insensitive
2 flow-sensitive context-sensitive invocation graph. By path-insensitive we mean that all possible
3 instruction execution paths within each method are included in the program graph. The program
4 graph is flow-sensitive since we consider the order of execution of the instructions in each
5 method, accounting for local variable “kills” and casting of object references. By
6 context-sensitive we mean that for each call site in a method, we identify the set of receivers,
7 target methods and parameters. Each node in the program graph includes this context, or state
8 information. This context is used in identifying a call site’s successor node in the program graph.

9 In addition, an advantageous embodiment of this invention includes a data flow analysis based
10 on the program graph. This data flow analysis is used to propagate object allocation sites and
11 string constants through the program graph. In some embodiments it is sufficient to only
12 propagate a limited number of allocation sites, including Permission, and Thread objects,
13 including all of their subtypes, as well as string constants. In a specific embodiment, , there is
14 primary interest in string constants used as parameters to Permission (or subtype) object
15 constructors. These Permission (or subtype) objects are subsequently passed to
16 AccessController.checkPermission() calls. As a practical matter, most parameters to Permission
17 (and their subtype) constructors are string constants. In a few cases the parameter value is the
18 result of a computation (e.g., string concatenations). A more sophisticated data flow analysis is
19 used in such circumstances.

20 For a “closed-world” analysis (e.g., the analysis of a whole program) the root nodes of the graph
21 are the entry points to the program. As an example, the root node for a stand-alone Java program
22 has a method signature of *static void main(String[])*. For applets, servlets, etc., multiple entry
23 points are considered as root nodes. For “open-world” analysis, each of the methods accessible
24 to applications using the library (e.g., *public* and *protected* methods) is included as a root node in
25 the graph.

1 An advantageous embodiment of the current invention employs a program graph with the
2 following characteristics. It is an invocation graph, where each node in the program graph
3 represents a method invocation. Each node in the graph includes the following context state:

- 4 • The target method.
- 5 • For instance methods, the type or allocation site of the method's receiver. The receiver
6 may include more detailed information, such as a string constant or memory allocation
7 site for the receiver. For static methods, there is no receiver.
- 8 • All parameters to the method, represented as a vector of sets of possible allocation sites
9 (or types) or constants. As with the receiver, it may include more detailed information.
- 10 • A set of possible return value allocation sites, constants or types from this method. As
11 with the receiver, it may include more detailed information.

12 The edges in the call graph represent call sites in methods to target methods. The edges in the
13 graph are directed, where each edge points from a calling method (call site) to a target method.
14 The graph may be cyclic. Also, each node in the graph is unique – its context (target method,
15 receiver, parameters vector) is unique. That is, two nodes in the graph will not have the same
16 context.

17 For simplicity of presentation, we assume bi-directional traversal of the program graph is
18 possible, even though the edges in the graph are nominally unidirectional. That is, from a root
19 node, we can traverse the entire program graph. Also, from any node within the program graph,
20 we can find all of its predecessor nodes.

21 **The Basic Automated Program Resource Identification and Association Algorithm**

22 The objective of the algorithm is to statically identify a set of authorizations used for each class
23 being analyzed. Specifically, the algorithm identifies methods (and their declaring classes) in
24 paths that pass through nodes representing calls to an authorization test. In Java 2 that
25 authorization method is `AccessController.checkPermission()` or a call to some of the

1 java.lang.SecurityManager authorization testing methods. Root nodes in the program graph are
2 included in the set of *start nodes*. Also, in Java 2, any AccessController.doPrivileged() node's
3 predecessor nodes are considered *start nodes*. A *stop node* is a node that is an authorization test.
4 In Java 2 a stop node is a node whose target method is AccessController.checkPermission() or
5 one of the SecurityManager authorization testing methods.

6 In an advantageous embodiment written for Java programs, the algorithm identifies all nodes in
7 all paths bounded by any *start node* and a *stop node* and associates a java.security.Permission *p*
8 (or an appropriate subtype) with each of the nodes in the path. The value of *p* is obtained by
9 identifying the parameter *p* to the authorization test *stop node*. The object *p* is an allocation site
10 obtained as a result of the data flow analysis. We take all of the methods that are represented by
11 the nodes in the paths, and identify the classes that declared the methods. Each such class is
12 associated with Permission *p*.

13 The following example of pseudo-code embodies a basic algorithm of the present invention:

```
14 // Identify all start and stop nodes in the graph  
15  
16 // The start set includes all nodes in the call graph root set  
17 Set start = new Set( rootNodes );  
18  
19 // Initially, the stop set is empty  
20 Set stop = new Set( );  
21  
22 // Identify additional start nodes and the stop nodes  
23 // by iterating over all nodes in the call graph.  
24 Iterator nodesIter = graphNodes.iterator( );  
25 while ( nodesIter.hasNext( ) ) {  
26     Node node = nodesIter.next( );  
27     if ( isDoPrivileged( node ) {  
28         start.addAll( node.getPredecessors( ) );  
29         continue;  
30     }
```

```

1      }
2      if (isCheckPermission( node ) ) {
3          stop.add( node );
4          continue;
5      }
6  }

7  // Find all nodes between stop to start nodes and get the used
8  // Permission. For each node, get its method and associated class.
9  // Associate the Permission with each class identified.

10 // For each checkPermission(perm) call, identify the Permission
11 // perm and the classes that need the Permission.
12 Iterator stopIter = stop.iterator( );
13 while ( stopIter.hasNext( ) ) {
14     Node stopNode = stopIter.next( );
15
16     // Get the Permission from the checkPermission() node
17     UsedPermission perm = getPermission( stopNode );
18
19     // Using a work list algorithm, find all nodes in all paths
20     // that are bounded by the nodes in the start set and
21     // the stop node. A work list algorithm is used since
22     // the graph may be cyclic.
23     Set nodes = getNodes( stopNode, start );
24
25     // For each such node, get the node's method and
26     // the class that declared that method.
27     // Add the Permission as being used for the class.
28     Iterator nodesIter = nodes.iterator( );
29     while ( nodesIter.hasNext( ) ) {
30         Node node = nodesIter.next( );
31         Method method = node.getMethod( );
32         Class declaringClass = method.getDeclaringClass( );
33         // Remember that this class needs this Permission
34         usedPerms.add( declaringClass, perm );
35     }
36 }

```

At the end of this algorithm, each class in the usedPerms map includes a set of Permissions used for the class.

Threads

When a Thread constructor is called, an AccessControlContext is instantiated and added to the new Thread instance as its inherited AccessControlContext. Subsequently, whenever the child Thread creates an AccessControlContext (e.g., during a call to checkPermission() or when a new Thread is created), the ProtectionDomains from the child Thread's inherited AccessControlContext's ProtectionDomains are included in the newly created AccessControlContext.

It should be noted that the creation of a new Thread does not result in the newly created Thread being started. In fact, the Thread could be started by a different Thread, or by a method in a class other than that which created the Thread. Nominally, the Thread.start() method calls the Thread.run() method. The run() method becomes the root (starting) method for the child Thread.

To correctly model a Thread, we should appropriately model its *constructor* and start() methods. We need to know which classes are included in a newly created Thread's inherited AccessControlContext. We can simplify the problem definition by rewriting the invocation graph to include a predecessor edge from a Thread.run() call site to the allocation site for the inherited AccessControlContext. The AccessControlContext allocation site is associated with the allocation site for the run() method's receiver (a Thread object). In general, an allocating method also calls an object's constructor.

1 To correctly model the inherited AccessControlContext allocation sites, we build a lookup table
2 that maps Thread allocation sites to the graph nodes where the respective inherited
3 AccessControlContext constructor is called. This mapping allows us create the new predecessor
4 edge for the Thread.run() method.

5 The following pseudo-code embodies the basic algorithm.

```
6 // For all Thread allocation sites, build a table that
7 // maps the Thread to its inherited AccessControlContext
8 // constructor.
9 Map threadInheritedACCMMap = new Map( );

10 // Iterate over all of the object allocations, selecting
11 // Thread allocations. From the Thread allocation site, obtain
12 // the inherited AccessControlContext constructor node.
13 Iterator allocIter = allocationSites.iterator( );
14 while ( allocIter.hasNext( ) ) {
15     AllocSite allocSite = allocIter.next( );
16     if ( allocSite.getClass( ) == java.lang.Thread ) {
17         AllocSite inheritedACC = allocSite.getInheritedACC( );
18         Node node = inheritedACC.getNode( );
19         threadInheritedACCMMap.add( allocSite, node );
20     }
21 }
```

22 Now, when we reach a Thread.run() node in the invocation graph, we can find its predecessors
23 by looking up the Thread allocation site's inherited AccessControlContext allocation site node
24 and use it as the replacement predecessor edge.

25 From the algorithm above, the getNodes() method is suitably modified to use allocCallSites
26 to find replacement predecessor nodes when searching the call graph.

doPrivileged with an AccessControlContext

In addition to the `AccessController.doPrivileged()` method described above that takes a `PrivilegedAction` argument, the other form of the `doPrivileged()` method also takes an `AccessControlContext` instance as an argument. The behavior of this method is similar to the version of the `doPrivileged()` that was previously described. However, in this case, when a `AccessController.checkPermission()` authorization test is performed, all of the `ProtectionDomains` in the `AccessControlContext` argument are also included in the authorization test.

To model this behavior, we want to include all of the classes that would be considered when the `AccessControlContext` was created. These classes will use the `Permission p` passed as an argument to `AccessController.checkPermission()`. To do so, we add the node for the allocation site for the `AccessControlContext` as a predecessor of the `doPrivileged()` call site. The `getNodes()` method in the algorithm will then correctly include all of the `AccessControlContext` nodes.

The precision of the invocation graph and data flow analysis affects the precision of the analysis as described by this invention. The compiler community commonly uses variants of one of several popular program graph construction techniques, such as Rapid Type Analysis (RTA) and Class Hierarchy Analysis (CHA). These algorithms are considered to be “conservative”. When writing an optimizer for a compiler, being “conservative” will not lead to incorrect execution of the resulting program, though the optimized code may run as fast as it could. However, these algorithms generate paths in the program graph that do not exist in the actual program. A graph that is overly conservative will result in the allocation of Permissions to classes that do not use the Permissions, possibly opening security holes. This would violate the “principle of least privilege”. Thus, it is desirable to employ a graph construction algorithm that minimizes its conservativeness.

1 To improve precision, an advantageous embodiment uses a program graph algorithm that is
2 context-sensitive and flow sensitive. Other graph construction algorithms can be employed, such
3 as those described above that are context-insensitive. However, people experienced in the art are
4 aware that a context-sensitive call graph yields more precise results.

5 A flow sensitive graph construction technique considers the order of execution of instructions
6 both intra and inter procedurally. A flow sensitive analysis improves the accuracy of the
7 resulting program graph construction. An advantageous embodiment uses a program graph that
8 is flow sensitive, including support for local field “kills” – local fields that are overridden by
9 subsequent assignments to the same field – and type casting.

10 Path insensitive static analysis of programs does not consider input values or the values of
11 constants defined within the program. All intra-procedural paths are considered during the
12 analysis. For example, in the statement:

13 *if (false) exp1 else exp2*

14 both *exp1* and *exp2* are evaluated even though, in practice, *exp1* may never get executed. While
15 conservatively correct, this results in additional nodes being generated in the graph that may not
16 occur in practice. The net effect is that all used Permissions are included, though some addition
17 Permissions may be included that are not strictly necessary in practice. Adding path sensitivity
18 would further reduce the conservativeness of the analysis.

19 In an advantageous embodiment, we define an allocation site to be the type of the object being
20 allocated (e.g., the class), the method in which the memory is being allocated, and the instruction
21 offset in that method. The problem that arises with this definition of an allocation site is that it
22 does not directly map to a node in the invocation graph. A change to the definition of the
23 allocation site that results in a one-to-one mapping will yield a more precise result.

1 In an advantageous embodiment of the present invention we are able to determine, for a given
2 application or classes in a library, the set of Java authorizations for each class within the
3 analysis scope. By using a context-sensitive flow-sensitive call graph, we are able to accurately
4 identify the classes in each path that include a call to the Java 2 security authorization subsystem.
5 The level of precision we attain is far greater than that used for Java 2 security since we are able
6 to identify authorization use to the level of instructions, basic blocks and methods, rather than
7 the coarser granularity of classes or *ProtectionDomains* (which may include many classes).

8 By using the analysis technique described in this invention, we can determine the Permissions
9 needed by mobile code, such as applets, servlets, and Enterprise Java Beans. Prior to loading the
10 code, it is possible to prompt the administrator or end-user to authorize or deny the code access
11 to restricted resources protected by the authorization subsystem.

12 This ability, to be able to automate the process of determining which authorizations need to be
13 granted to the application, changes the relationship between the developer of the code and the
14 administrator / end-user. Instead of relying solely on recommendations from the
15 developer, or resorting to trial-and-error testing of the code to determine required Permissions, a
16 tool can analyze the code and make its own recommendations and/or validate recommendations
17 made by the developer. This shifts the relationship from one that *requires* that the developer be
18 trusted, to something that can be verified.

19 Java 2 allocates Permissions to classes. As a practical matter, Permissions are aggregated at the
20 level of *ProtectionDomains* (e.g., at the JAR or file system directory level). This coarsens the
21 granularity when assigning Permissions. Without the present invention, this level of granularity
22 appears to be reasonable. The present invention allows identification of authorization usage
23 down to a specific method. A refinement of the Java 2 authorization algorithm could result in
24 the minimization of authorization, getting us closer to the application of the “principle of least
25 privilege”. In addition, anyone skilled in the art could apply the present invention to other

languages, including but not limited to the Microsoft C# language, Including but not limited to procedural languages and functional languages.

Thus in an example embodiment of the present invention, as is shown in Figure 1, a method is provided including the steps of employing a computer for obtaining a collection of code {110}, providing a program graph representing the collection of code {120}, identifying authorization resources in the collection of code {130} locating bounded paths within the program graph {140} and associating the located authorization resources with the bounded paths {150}. When necessary, perform a data flow analysis (160) to help in the identification of the resource on which authorization testing is to be performed. The data flow is used to propagate constants and/or symbolic representations of the program graph which are used for identifying state information that could be passed to a authorization test. In its simplest form, a constant is passed as a parameter to an authorization test.

In some embodiments of the method, the collection of code includes codes obtained from a group of codes including basic blocks, class methods, classes, collections of classes, object code, or any combination of these; and/or the step of providing a program graph 120, includes constructing the program graph through static analysis techniques.

The identifying authorization resources may include finding at least one authorization point in the program graph. and/or a Java authorization point. Finding the Java authorization point may also include finding an AccessController.checkPermission node in the program graph, and finding a java.security.Permission object passed as an argument to the AccessController.checkPermission method.

In other embodiments, the authorization point is an instruction invocation. The instruction may be a call to an operating system or supervisor program, where the instruction (or a numeric constant parameter) indicates the resource being accessed. In still other embodiments, the

1 instruction invocation is used in a particular language for the collection of code, which includes
2 but is not limited to C, C++, and C#.

3 The step of identifying authorizations 130 may include employing data flow analysis. The data
4 flow analysis is often used to propagate constants or symbolic references to values which further
5 discriminate the resource for which authorization is being performed. In such cases, the process
6 will need the set of possible values each variable, parameter or receiver could hold at various
7 points in the program.

8 Identifying or locating any bounded path 140 may also include locating a set of start nodes and
9 locating a stop node in the program graph. A bounded path includes all nodes within the graph
10 bound by start nodes and a stop node.

11 The step of associating authorizations with the graph may also include associating and
12 aggregating the any authorization resource with the collection of code. This step may include
13 associating the authorization resource with the basic block, class method, method and/or
14 collection of these forms of collection of code. The granularity of the aggregation varies
15 depending on the need of the authorization system or system that uses the results of the algorithm
16 described in the present invention.

17 In an example embodiment of the present invention, the invention includes an apparatus for a
18 computing means 200 providing advantageous results. As shown in Figure 2, the computing
19 means 200, includes means for obtaining a collection of code (210), means for providing a
20 program graph (220) representing the collection of code, means for identifying authorization
21 resources (230) used by the collection of code, means for locating a bounded path (240) within
22 the program graph, and means for associating an authorization resource (250) with the bounded
23 path. The components of the computing means are generally collocated. In alternate

embodiments the components are not collocated but are coupled to communicate and act together to perform the function of the apparatus.

In such an embodiment, the collection of code includes, but is not limited to, codes obtained from a group of codes including basic blocks, class methods, classes, collections of classes or any combination of these. The program graph from the collection of code can be obtained through static analysis techniques. The static analysis can be performed on object code, source code, an intermediate representation of the program, such as is generated by a compiler, or any other suitable form of the code, or any suitable combination thereof.

The means for finding an authorization resource includes locating at least one authorization point in the program graph. For example, when the program is written in Java, the program graph is searched for one of the Java authorization points -- a call to `AccessController.checkPermission()` or one of the authorization test methods in the `java.lang.SecurityManager` class. In the case of calling `AccessController.checkPermission()`, the means for performing data flow analysis (260) will propagate an object which represents the authorization resource. In Java, this object is assignment compatible with `java.security.Permission`, and the value passed as a parameter to the `AccessController.checkPermission()` method. Alternatively, in other programming languages or systems, the authorization resource is a specific instruction, such as an operating system or supervisor call, or a set of specific method or procedure calls. Other languages may include, but are not limited to, C, C++ and C#.

Some authorization mechanisms are in effect over a bounded range of code within the program graph. In this case, the embodiment of the present invention provides means for locating a bounded path (240) includes locating a set of start nodes in the program graph, and locating a stop node in the program graph, where the bounded path includes all nodes within the graph bound by the start nodes and said stop node. In addition, the means for associating authorization resources with a bounded path (250) was previously described.

1 In an example embodiment of the present invention, the invention includes a method employing
2 a computer 300 which enables the identification of a complete set of authorization resources of
3 the collection of code. As shown in Figure 3, the method includes the steps of obtaining a
4 collection of code (310), providing a program graph (320) representing the collection of code,
5 and identifying a complete set of authorization resources (330) of the collection of code. If no
6 authorization resources are identified in step (330), this indicates that authorization testing is not
7 necessary for the collection of code. In the case of the runtime environment being Java, and
8 when no other collections of code other than those identified in step (310) are running in the
9 same runtime environment, then it may be possible for the runtime environment to run without a
10 java.lang.SecurityManager installed, or to perform any other authorization tests during the
11 execution of the collection of code.

12 It is noted that heretofore there was no method which was able to guarantee that it identified a
13 complete set of authorization resources (330) of the collection of code. This is because existing
14 program graph construction techniques did not scale well to large collections of code; and the
15 analysis of non-type safe languages led to incomplete results due to the lack of precision in
16 existing points-to analysis techniques. In addition, other existing type safe languages have not
17 included security features that would require a search of authorization resources. It is very
18 advantageous to identify the complete set of authorization resources of the collection of code in
19 order to ensure that the correct set of authorization resources can be granted to the collection of
20 code to allow it to execute. The alternative is that an insufficient set of authorizations are
21 granted, thus prohibiting the program from executing; or too many authorizations are granted,
22 resulting in possible security holes. Another reason for searching for the authorization resources
23 is to identify possible resources which may be accessed, but to which a system administrator
24 (e.g., a web browser user) would like to prohibit the collection of code from accessing.

1 In still another example embodiment, the invention includes an apparatus for a computer
2 providing the advantages of the present invention. As shown in Figure 4, a computer (410) with
3 access to a collection of code, includes an authorization resource identifier (430) to identify any
4 authorization resources within the collection of code, a bounded path locator (440) to locate
5 bounded paths within a program graph of the collection of code, and an associator (450) to
6 associate any authorization resources with located bounded paths. The components of the
7 apparatus are generally collocated. In alternate embodiments the components are not collocated
8 but are coupled to communicate and act together to perform the function of the apparatus.

9 The collection of code generally includes codes obtained from a group of codes including basic
10 blocks, class methods, classes, collections of classes or any combination of these. If a program
11 graph exists for this collection of code, the program graph is acted upon by the bounding path
12 locator. When a program graph does not exist, the collection of code is passed to a program
13 graph constructor (420) which produces a program graph. It is advantageous for the program
14 graph constructor (420) to employ static analysis techniques in producing the program graph.

15 The static analysis can be performed on object code, source code, an intermediate representation
16 of the program, such as is generated by a compiler, or any other suitable form of the code, or any
17 suitable combination thereof.

18 As previously described, the authorization resource identifier (430) finds an authorization
19 resource. If no authorization resources are identified, this indicates that authorization testing is
20 not necessary for the collection of code. In the case of the runtime environment being Java, and
21 when no other collections of code other than those identified by the computer (410) are running
22 in the same runtime environment, then it may be possible for the runtime environment to run
23 without a java.lang.SecurityManager installed, or to perform any other authorization tests during
24 the execution of the collection of code. When the program is written in Java, the program graph
25 is searched for one of the Java authorization points -- a call to
26 AccessController.checkPermission() or one of the authorization test methods in the

1 java.lang.SecurityManager class. In the case of calling AccessController.checkPermission(), the
2 means for performing data flow analysis (460) will propagate an object which represents the
3 authorization resource. In Java, this object is assignment compatible with
4 java.security.Permission, and the value passed as a parameter to the
5 AccessController.checkPermission() method. Alternatively, in other programming languages or
6 systems, the authorization resource may be a specific instruction, such as an operating system or
7 supervisor call, or a set of specific method or procedure calls. Other languages may include, but
8 are not limited to, C, C++ and C#. In some cases, the identification step (430) employs a data
9 flow (460). This data flow is often generated from the program graph.

10 In still another example embodiment of the present invention, the invention includes an
11 apparatus for a computing module (510), shown in Figure 5. The computing module (510) has
12 access to a collection of code, and in some cases to a program graph for the collection of code.
13 The computing module (510) includes: an authorization resource identifier (530) to identify any
14 authorization resources within the collection of code; a program graph constructor (520) to
15 construct a program graph when a program graph is not available; and a bounded path locator
16 (540), to locate bounded paths within the program graph of the collection of code. If no
17 authorization resource is identified by the authorization resource identifier (530), the identifier
18 provides an indication that authorization testing is not necessary. It is advantageous for the
19 program graph constructor (520) to use static analysis techniques to construct the program graph
20 when a program graph is not available. The components of the computing module are generally
21 collocated. In alternate embodiments the components are not collocated but are coupled to
22 communicate and act together to perform the function of the apparatus.

23 In another example embodiment of the present invention, the invention includes an apparatus for
24 identifying a complete set of authorization resources of a collection of code, as is shown in
25 Figure 6. The apparatus includes computing means 600. The computing means includes: means
26 for obtaining a collection of code (610); means for providing a program graph representing said

1 collection of code (620); and means for identifying a complete set of authorization resources
2 (630) of the collection of code. In some cases there is also means for indicating whether
3 authorization testing is or is not necessary (640). If no resource is identified in step (630), this is
4 an indication that authorization testing is not necessary.

5 In a further example embodiment of the present invention, the invention includes a method
6 employing a computer with access to a collection of code (710), shown in Figure 7. The
7 method includes the steps of obtaining or constructing a program graph (720) from the
8 collection of code [using static analysis techniques]; performing a data flow analysis (760) [using
9 static analysis techniques]; searching the program graph for useful resource (730) for executing
10 the collection of code; identifying any bounded paths (740) within the program graph over
11 which the useful resource is beneficial; and associating the useful resource (750) with the
12 collection of code. When used, the static analyses can be performed on object code, source code,
13 an intermediate representation of the program, such as is generated by a compiler, or any other
14 suitable form of the code, or any suitable combination thereof. The program graph can be a call
15 graph, an invocation graph, and/or any other suitable graph as is known to those skilled in the art.
16 The program graph can be context sensitive, which may include type information for any method
17 receiver and/or any of the parameters. As is known to those skilled in the art, the type
18 information may be further refined to include class and memory allocation site information
19 and/or any other differentiator which would further refine the labeling of the type, including per
20 instance information. The collection of code includes codes constructed from a group of codes
21 including basic blocks, class methods, classes, collections of classes or any combination of these.

22 In some embodiments, the step of searching (730) includes locating a node or edge in the
23 program graph (720) that represents a location where the useful resource (730) would be
24 beneficial. This location may be an authorization test.

1 In further embodiments, the program graph (720) represents an object oriented program, such as
2 a program written in the Java programming language. The program graph (720) and data flow
3 analysis (760) are obtained from Java object code. The useful resource identified in step (730)
4 may be a resource identifier, such as a java.security.Permission object or an object that is
5 assignment compatible. The authorization test may be a call to any

6 `java.security.AccessController.checkPermission`

7 method or any authorization testing method in an instance of

8 `java.lang.SecurityManager`

9 and/or one of its subclasses. One of these authorization tests may accept, as a parameter, a

10 `java.security.Permission`

11 object or an object that is assignment compatible. Through the use of the program graph and
12 data flow analysis, the constructor of the

13 `java.security.Permission`

14 object (or subclass) is located, and the data flow analysis is used to identifying any value passed
15 by as a parameter to the constructor, where the combination of the Java

16 `java.security.Permission`

17 and a value for any parameter used permission for the authorization test.

18 In addition, the step of identifying bounded paths (740) often includes locating a set of start
19 nodes in the program graph, and locating a stop node in the program graph, and the bounded
20 paths include all nodes within the program graph bound by the start nodes and stop node.

21 In addition, in the step of associating (750), the useful resource is sometimes mapped using any
22 number of techniques known to those skilled in the art. This mapping may include the use of
23 hash tables that use an element of the collection of code as the key and a value that holds a
24 collection of java.security.Permission objects. Thus it is possible to use a subset of the collection
25 of code to the identified useful resource used by that subset of the collection of code. This subset
26 may include the nodes in the bounded path.

1 In an example embodiment of the case of code written to use the reference implementation of the
2 Java runtime environment, the stop node represents the method

3 `java.security.AccessController.checkPermission(),`

4 and the start nodes are any of root nodes in the program graph or a node representing the method

5 `java.security.AccessController.doPrivileged().`

6 In this case, a

7 `java.security.Permission`

8 object, or an object that is assignment compatible, is associated with the code associated with
9 the bounded path. In addition, the same

10 `java.security.Permission`

11 object, or an object that is assignment compatible, is associated with the code associated with any
12 node in the program graph prior to the

13 `java.security.AccessController.doPrivileged()`

14 node. Also, a

15 `java.security.Permission`

16 object, or an object that is assignment compatible, is associated with each

17 `java.security.AccessController.checkPermission()`

18 in the program graph.

19 As previously described, after Permission objects are associated with nodes in the program
20 graph, generally, the Permissions are then associated with method that is represented by with the
21 node in the program graph. As previously described, all Permissions associated with each
22 method are also associated with their respective declaring classes. As previously described, if
23 the collection of code (710) is a set of classes, then the Permissions associated with the classes in
24 the set also are associated with the set of classes. That is, the Permissions are aggregated from
25 the smaller collections of code to the larger collections of code. It is advantageous that during

1 the execution of the collection of code (710), the useful resources identified in the present
2 invention is employed when executing the collection of code.

3 In still a further example embodiment of the present invention, the invention includes a method
4 for statically detecting useful resources for a collection of code 800, shown in Figure 8. The
5 code is written in a computer programming language. As shown in Figure 8, the method
6 includes the steps of: calculating if any code in a collection of code is part of a program graph
7 (810); identifying any useful resource (820) in the collection of code; determining bounded paths
8 of nodes within the program graph which constrain the useful resources (830) of the collection of
9 code; and associating the useful resource with the collection of code within the bounded path
10 (840) of nodes in the program graph.

11 In useful embodiments of the present invention, the program graph can be a call graph, an
12 invocation graph, and/or other suitable graph known to those skilled in the art, which have a set
13 of root nodes. From the program graph it is possible to determine whether a basic block,
14 method, class is represented as a node in the graph. By using the program graph, it is possible to
15 identify used resources. The bounded nodes are all nodes in the graph that are bounded between
16 start nodes and a stop node, and all nodes in between the start nodes and the stop node. The start
17 nodes may include the program graph root nodes.

18 Thus in particular embodiments of a method statically detecting useful resources for a collection
19 of code: the step of calculating includes using said program graph is an invocation graph includes
20 having a set of root nodes for said program graph; and/or the step of calculating includes using a
21 call graph having a set of root nodes for said program graph; and/or the method includes: a step
22 of determining whether a basic block, method, class is represented as a node in said invocation
23 graph; a step determining whether a basic block, method, class is represented as a node in said
24 call graph; and/or a step identifying a node in said invocation graph that identifies said useful
25 resources, wherein said useful resources is one or more resources; and/or a step identifying a

1 node in said invocation graph that identifies said useful resources, wherein said useful said
2 resources is one or more resources; and/or the nodes in said bounded path of nodes bounded
3 nodes are all nodes in the graph that are bounded between start nodes and a stop node; and/or the
4 nodes in the bounded path of nodes are all nodes in the graph that are bounded between a set of
5 start nodes and a stop node; and/or some or all root nodes in said invocation graph are members
6 of the start nodes.

7 It is noted that the present invention can be realized in hardware, software, or a combination of
8 hardware and software. A visualization tool according to the present invention can be realized in
9 a centralized fashion in one computer system, or in a distributed fashion where different elements
10 are spread across several interconnected computer systems. Any kind of computer system - or
11 other apparatus adapted for carrying out the methods described herein - is suitable. A typical
12 combination of hardware and software could be a general purpose computer system with a
13 computer program that, when being loaded and executed, controls the computer system such that
14 it carries out the methods described herein. The present invention can also be embedded in a
15 computer program product, which includes all the features enabling the implementation of the
16 methods described herein, and which - when loaded in a computer system - is able to carry out
17 these methods.

18 Computer program means or computer program in the present context include any expression, in
19 any language, code or notation, of a set of instructions intended to cause a system having an
20 information processing capability to perform a particular function either directly or after either
21 or both of the following a) conversion to another language, code or notation; b) reproduction in a
22 different material form.

23 It is noted that the foregoing has outlined some of the more pertinent objects and embodiments of
24 the present invention. This invention may be used for many applications, such as gathering of
25 information to determine authorizations used by collections of code. Thus, although the

1 description is made for particular arrangements and methods, the intent and concept of the
2 invention is suitable and applicable to other arrangements and applications. It will be clear to
3 those skilled in the art that modifications to the disclosed embodiments can be effected without
4 departing from the spirit and scope of the invention. The described embodiments ought to be
5 construed to be merely illustrative of some of the more prominent features and applications of
6 the invention. Other beneficial results can be realized by applying the disclosed invention in a
7 different manner or modifying the invention in ways known to those familiar with the art.

COPIES OF THIS DOCUMENT ARE AVAILABLE FROM THE NATIONAL ARCHIVES AT COLLEGE PARK, MARYLAND